

STORING FIBRE CHANNEL INFORMATION ON AN  
INFINIBAND ADMINISTRATION DATA BASE

BACKGROUND OF THE INVENTION

FIELD OF THE INVENTION

[0001] This invention relates to an Infiniband™ (IB) input/output unit (IOU) that enables a host computing system to access Input/Output (I/O) devices that conform to the Small Computer System Interface (SCSI) mapping onto Fibre Channel (FC). Such I/O devices are referred to as FCP I/O devices.

DESCRIPTION OF BACKGROUND

[0002] The following references pertain to this invention:

1. Information regarding IB subnet management and administration can be found in Chapters 14, 15 and 16 of Infiniband Architecture Specification, Volume 1, release 1.1, Infiniband Trade Association. Infiniband is a trademark of SYSTEM I/O doing business as InfiniBand Trade Association of Portland Oregon.
2. Information regarding Fibre Channel can be found in Fibre Channel-Framing and Signalling (FC-FS) rev 1.9, American National Standards Institute, Inc.
3. Information regarding FCP I/O devices can be found in Fibre Channel Protocol for SCSI, second version, (FCP-2), American National Standards Institute, Inc.
4. Information regarding the SCSI mapping onto IB can be found in SCSI Remote Direct Memory Access (RDMA) Protocol (SRP), rev 16a, American National Standards Institute, Inc. and
5. Information regarding SRP target port identifiers can be found in Fibre Channel HBA API (FC-HBA), rev 8, American National Standards Institute, Inc.

[0003] In an Infiniband configuration containing an I/O Unit (IOU) that enables the host computing system connected to an Infiniband network to access I/O devices, the FCP I/O devices appear to the host as SRP I/O devices. Since all SRP I/O devices are uniquely identified by SRP target port identifiers (IDs), there is a one-to-one correspondence between the worldwide-unique port name (WWPN) for each FCP I/O device and a corresponding SRP target port identifier. SRP target port identifiers consist of two 64-bit quantities. The first is an “IOCGUID” that identifies the IOC that controls the SRP target port. The second is a 64-bit extension. For the case of an IB-to-FC adapter, the extension is set to the WWPN of the FCP I/O device corresponding to the SRP target port identifier. For example, the SRP target port identifier IOCGUID1.WWPNA corresponds to an FCP I/O device with WWPN equal to WWPNA that is accessible through an IOC with an IOCGUID equal to IOCGUID1; SRP target port identifier IOCGUID1.WWPNB corresponds to an FCP I/O device with WWPN equal to WWPNB this is also accessible through an IOC with IOCGUID equal to IOCGUID1; SRP target port identifier IOCGUID2.WWPNC corresponds to an FCP I/O device with WWPN equal to WWPNC that is accessible through an IOC with an IOCGUID equal to IOCGUID2; and so on.

[0004] In addition to a target port identifier, SRP I/O devices are identified by an IB “service name.” IB service names uniquely identify “service providers,” such as information databases, communication functions, or (as in this case) SRP I/O device functions. The SRP I/O device service name is a 40-byte UTF-8 character string. The first 24-bytes of the service name contains the following character string:

‘SRP.T10:xxxxxxxxxxxxxxxx’

[0005] The ‘SRP.T10’ part of the service name identifies the service as an SRP target port; the ‘xxxxxxxxxxxxxxxx’ part of the service name is a 16-byte hexadecimal encoding of the 64-bit “extension” portion of the SRP target port identifier. Note that in the case of IB-to-FC adapters, the 64-bit extension of the SRP target port identifier is set to the WWPN of the FCP I/O device corresponding to the SRP target port; therefore, the FCP I/O device corresponding to a particular service name can be deduced by examination of the ‘xxxxxxxxxxxxxxxx’ portion of the service name. The additional 16 bytes of the SRP service name contain null characters. An

example of an SRP service name corresponding to an FCP I/O device with WWPN x'5347 9899 5348 8888' is 'SRP.T10:5347989953488888', followed by null characters.

[0006] Since the service name and SRP target port identifier formats described above contain the FC WWPN of the corresponding FCP I/O device, a host is able to generate the SRP service name for the SRP I/O device that corresponds to the FCP I/O device. Since hosts typically identify FCP I/O devices by their WWPNs, this enables an easy determination of the SRP service name for the corresponding SRP I/O device whenever the host needs to access a particular FCP I/O device with a given WWPN.

[0007] In order to access the SRP I/O device, however, the host also needs the IB address of the IB-to-FC adapter supporting the service. Prior to this invention, there was no efficient method by which a host could determine the IB address of the IB-to-FC adapter that supports the service name corresponding to the FCP I/O device. Therefore, without the use of this invention, a host attempting to access an FCP I/O device with a particular WWPN would need to poll all of the I/O units (IOUs) in the IB subnet in order to determine those IOUs which are IB-to-FC adapters. The host would then poll each IB-to-FC adapter to determine the particular IB-to-FC adapter (or adapters) that provide access to the FCP I/O device with the desired WWPN. Such a polling operation, which involves sending multiple queries to each IOU in the subnet, is impractical when the IB subnet is of any significant size, and results in unacceptable performance degradation.

## SUMMARY OF THE INVENTION

[0008] This invention defines the means for efficiently storing and retrieving information about FCP I/O devices on an IB Subnet Administration Database. The method by which information is stored enables a host to rapidly determine the IB addressing parameters by which FCP I/O devices are accessed through the IB-to-FC adapter.

[0009] The shortcomings of the prior art, namely the lack of an efficient means to discover the IB address of the IB-to-FC adapter providing access to a particular FCP I/O device identified by a WWPN, are overcome by the registration of SRP service names on an IB subnet

administration database, and by the querying of that database as described in this invention. Use of this invention decreases the complexity of host systems, improves their performance, and provides additional advantages as outlined below.

[0010] Additional features and advantages are realized through the techniques of the present invention. Other embodiments and aspects of the invention are described in detail herein and are considered a part of the claimed invention. For a better understanding of the invention with advantages and features, refer to the description and to the drawings.

#### BRIEF DESCRIPTION OF THE DRAWINGS

[0011] The foregoing and other objects, features, and advantages of the invention are apparent from the following detailed description taken in conjunction with the accompanying drawings in which:

[0012] FIG. 1 illustrates one example of a network configuration containing an IB subnet, an FC fabric, and an IB-to-FC adapter which incorporates this invention;

[0013] FIG. 2 illustrates one example of an IB-to-FC adapter and associated components which incorporates this invention;

[0014] FIG. 3 illustrates one example of an IB Service Record; and

[0015] FIG. 4 illustrates one example of a flow chart of the process of determining the IB address of the IB-to-FC adapter(s) providing access to a particular FCP I/O device.

[0016] The detailed description explains the preferred embodiments of the invention, together with advantages and features, by way of example with reference to the drawings.

#### DETAILED DESCRIPTION OF THE INVENTION

[0017] This invention provides host computing systems with the ability to efficiently identify the Infiniband (IB) addressing parameters which correspond to Fibre Channel I/O devices that are accessible through IB-to-Fibre Channel adapters. The invention can be subdivided into a

configuration step and a look-up step. In the configuration step, the FCP I/O devices are registered in a database. In the look-up step, a host accesses the database in order to determine the IB addressing parameters of the IB-to-FC adapter that provides access to the FCP I/O device. These two steps are described in detail below.

[0018] Referring to Fig 1, host systems 101a and 101b communicate with an IB-to-FC-adapter 105 through an IB subnet 104. The IB-to-FC adapter 105 communicates with the FCP I/O devices 107a – 107d through the FC fabric 106. Each host system, for instance host system 101a, includes a processing unit 110 for executing data processions instructions, and a memory 111 for storing data to be processed and encoded processing instructions arranged in computer programs forming commands, requests and routines. The interface protocol between the host systems 101 and the IB-to-FC adapter 105 conforms to ref. 4. Each FCP I/O device 107a – 107d appears to the host 101 as if it were an IB I/O device conforming to SRP (ref 4); therefore, the I/O devices 107 visible to the host 101 are referred to as SRP I/O devices even though they represent actual FCP I/O devices. The interface protocol between the IB-to-FC adapter 105 and the FCP I/O devices 107 conforms to the SCSI mapping onto FC (ref 3.).

[0019] Fig. 2 shows an expanded representation of an IB-to-FC adapter 105 containing two I/O controllers (IOCs), 201a and 201b. The IOCs 201 are attached to the FC fabric 106, which is attached to the FCP I/O devices 107a-107d. IOC 201a provides access to FCP I/O devices 107a and 107b, and IOC 201b provides access to FCP I/O devices 107c and 107d. All FCP I/O devices 107 are uniquely identified by a 64-bit “worldwide-unique” port name (WWPN). Thus, FCP I/O devices 107a – 107d are uniquely identified by WWPN A through WWPN D.

[0020] Returning to Fig. 1, during a configuration step, each FCP I/O device 107a -107d is registered with the subnet administration (SA) database 103. This registration can be done by the IB-to-FC adapter 105, or it might be done by a third party such as another host which executes a configuration routine.

[0021] In order to register an FCP I/O device 107 with the SA database 103, it is necessary to store, in the SA database 103, an IB service record which identifies the FCP I/O device 107. This

is accomplished by sending an IB datagram containing the SubnetAdminSet (ServiceRecord) command to the subnet manager 102. This command contains an IB service record, as shown in Fig. 3. The service record is stored in the subnet administration database 103. Referring to Fig. 3, the fields in the IB ServiceRecord which relate to this invention are the ServiceName 301 and ServiceData 302 fields. (The ServiceID, ServiceGID, ServiceP\_Key, ServiceLease, and ServiceKey fields are used as specified in IB specifications, ref 1.)

[0022] The first 24 bytes of the 512 bit (64-byte) ServiceName field 301 are set to the first 24 bytes of the SRP service name corresponding to the FCP I/O device. These bytes are followed by the string ‘.FCP’ followed by a sequence of null characters to fill the remaining bytes in the ServiceName field. The string ‘.FCP’ is appended to the SRP service name to distinguish it from an SRP service name for an SRP I/O device that does not correspond to an FCP I/O device. Such a service name would contain null characters in the corresponding character positions.

[0023] The first 64-bits of the ServiceData field 302 are set to the IOCGUID of the IOC that provides access to the FCP I/O device. The remaining bytes in the ServiceData field are not used by this invention and may be set to any value. This completes the configuration step for the FCP I/O device; the configuration step is repeated for all FCP I/O devices accessible from the IB subnet.

[0024] Provided the above configuration steps have been completed, the host 101 is now able to determine the IB address of the IB-to-FC adapter 105 providing access to an FCP I/O device with a given WWPN by performing the steps shown in Fig 4.

[0025] Referring to Fig. 4, the host sends a SubnetAdminGet (ServiceRecord) request (401) to the subnet manager 102 (SM). The request indicates that the SM 102 is to return all service records containing the service name corresponding to the service name registered for the FCP I/O device 107 during the configuration step. For example, to obtain the service records corresponding to the FCP I/O device with a WWPN of x’5347 9899 5348 8888’, the host requests all service records containing a ServiceName field of

‘SRP.T10:5347989953488888.FCP’, followed by null characters. The procedure for sending this request is described in IB specifications ref. 1.

[0026] If an IB-to-FC adapter 105 in the IB subnet 104 provides access to the FCP I/O device 107 identified by the SubnAdminGet (ServiceRecord) request, then the response contains at least one service record corresponding to the FCP I/O device; if the response does not contain any service records, then the device is not accessible (402), and the procedure terminates.

[0027] Assuming the response contained at least one service record, then for each service record, a host 101 determines the IB addressing parameters necessary to access a FCP I/O device 107 as follows:

1. In step 403, the host 101 determines the IB addressing parameters of the IB-to-FC adapter 105 by converting the ServiceGID field into an IB “path.” The procedure for doing this is described in IB specifications.
2. In step 404, the host 101 determines the IOC 201 providing access to the FCP I/O device 107 by extracting the 64-bit IOCGUID from the ServiceData field. Note that the IOCGUID was stored in the ServiceData field 302 during the configuration step for the FCP I/O device 107.
3. In step 405, the host 101 constructs an SRP target port identifier corresponding to the FCP I/O device 107 by concatenating the IOCGUID and the WWPN of the I/O device 107. That is, the SRP target port identifier is set to IOCGUID.WWPN.
4. In step 406, the host 101 accesses the I/O device 107 by performing the procedures outlined in SRP specifications, ref 4.

[0028] The above steps are repeated for each service record returned in response to the SubnAdminGet(ServiceRecord) request. Performing the sequence of steps for each service record results in the identification of all physical paths to the FCP I/O device 107 through the IB subnet 104 and FC Fabric 106. The physical paths may include access through different IB-to-FC adapters, different IOCs within each adapter, different ports on the FC fabric, and different routes through the fabric.

[0029] Note that the above process did not require the host to poll multiple IOUs in the subnet prior to accessing the I/O device, as would have been required prior to this invention. Additionally, the host does not need to poll all the IOCs within an IOU in order to determine the IOC supporting the service name corresponding to the FCP I/O device. Prior to this invention, the host needed to poll up to 256 IOCs in each IOU in order to determine the desired IOC. Instead of performing all the above polling operations—an unacceptably long process, a host is able to determine all the physical paths to the FCP I/O device by sending a single request to the subnet administration database.

[0030] The capabilities of the present invention can be implemented in software, firmware, hardware or some combination thereof.

[0031] As one example, one or more aspects of the present invention can be included in an article of manufacture (e.g., one or more computer program products) having, for instance, computer usable media. The media has embodied therein, for instance, computer readable program code means for providing and facilitating the capabilities of the present invention. The article of manufacture can be included as a part of a computer system or sold separately.

[0032] Additionally, at least one program storage device readable by a machine, tangibly embodying at least one program of instructions executable by the machine to perform the capabilities of the present invention can be provided.

[0033] The flow diagrams depicted herein are just examples. There may be many variations to these diagrams or the steps (or operations) described therein without departing from the spirit of the invention. For instance, the steps may be performed in a differing order, or steps may be added, deleted or modified. All of these variations are considered a part of the claimed invention.

[0034] While the preferred embodiment to the invention has been described, it will be understood that those skilled in the art, both now and in the future, may make various improvements and enhancements which fall within the scope of the claims which follow. These claims should be construed to maintain the proper protection for the invention first described.